

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION

For

Improved Methods and Apparatus for Back-Projection and Forward-Projection

Assignee:

Mercury Computer Systems, Inc.

Inventors:

Scott A. Thieret, a U.S. citizen residing at 28
Nightingale Rd; Nashua, NH 03062

Joseph Goddard, a U.S. citizen residing at
72A Proctor St.; Salem MA

Ari P. Berman, a U.S. citizen residing at
60 North Street; Lexington MA

Valeri M. Kotlov, a U.S. citizen residing at
11 Saw Mill Brook Way; Woburn, MA

Background of the Invention

The invention relates to digital data processing and, particularly, to novel methods and apparatus for forward-projection and back-projection for use, e.g., in reconstructive imaging. The invention has application in health care and, specifically, medical imaging — as well as in a 5 host of other areas, such as manufacturing, security, research, and defense.

A *forward-projection* is an operation for estimating the (two-dimensional) transmission or emission shadow that would result from illumination of a (three-dimensional) object by a point source which emits radiation that partially passes through the object. By way of example, a 10 forward-projection could be performed in order to provide an estimate of what an x-ray image of a human might look like. The term *forward-projection* also refers to the estimate that is determined from executing a forward-projection operation.

A *back-projection* is an operation for generating a simple estimate of the structure of a (three-dimensional) object that could produce a given (two-dimensional) transmission or emission shadow, if that object was illuminated by a point source which emits radiation that at 15 least partially passes through the object. To continue the example, a back-projection could be performed in order to generate a simple estimate of the body structures the imaging of which resulted in a particular x-ray image. As above, the term *back-projection* also refers to the estimate that is determined from executing such a back-projection operation.

Reconstructive imaging is a procedure for constructing a three dimensional (3D) model of a 20 subject from a series of two-dimensional (2D) projection images. These projection images are typically generated by exposing the subject to a x-ray (or other radiation) point source that is positioned on one side of the subject and measuring the transmitted (or non-absorbed) radiation with a 2D detector that is positioned on the opposite side of the subject. A series of these projections is created by repeatedly measuring the transmitted radiation as the point source 25 and/or the detector are moved differing locations relative to the subject.

For example, computed axial tomography (CAT) is a reconstructive imaging technique in which a three-dimensional model of the internals of an object — typically, some portion of the human body — is constructed from a series of two-dimensional x-ray projections taken at evenly spaced angles that surround the object. With the aid of a computer, the series of acquired 2D 30 projection images are computationally projected backwards to estimate the relative density of the internal structures of the object (e.g., in the case of a head, the skull, gray matter and other

resolvable structures). While no one of these back-projections is, by itself, sufficient to identify those structures, the computational combination of all of the back-projections typically produces a reasonable 3D representation of the object. Because this typically involves performing back-projections on hundreds, if not thousands, of projected images, reconstruction systems typically 5 incorporate special-purpose hardware in order to provide timely results. Though prior art solutions have proven effective in this regard, they are typically quite expensive.

Moreover, those prior art solutions for determining back-projections have limited applicability to a new class of imaging, referred to as tomosynthesis, in which a relatively small 10 number of projection images are acquired from only a limited number of x-ray source positions (or foci). Typically, these foci are distributed across a limited arc. For this type of imaging, the requisite reconstruction requires an iterative approach. In these systems, the computations include not only back-projections, but also forward-projections, in which estimates of the volume 15 being reconstructed are projected "forward" to generate hypothetical projection images for computational comparison with actual measured projection images acquired by the imaging equipment. By comparing the hypothetical forward projection of the current 3D model to the measured projection images, a correction image can be calculated and used to update (modify) the current 3D model.

CAT scanners and tomosynthesis systems are not the only medical imaging equipment 20 that use forward-projections and/or back-projections. Forward-projection and back-projection operations also form the basis of a broader class of computer-based imaging techniques, referred 25 to as computed tomography (CT), as well positron emission tomography (PET), single photon emission computed tomography (SPECT), to name but a few. In addition, back-projection and forward-projection are used outside the medical field, for example, in manufacturing (e.g., to inspect articles for defects), security (such as baggage scanning) in research, defense and so forth.

An object of the invention is to provide improved digital data processing apparatus and methods. Another object is to provide such apparatus and methods as can be applied in reconstructive imaging applications, e.g., of the types referred to above (by way of non-limiting example).

30 Related objects of the invention are to provide such apparatus and methods as can be performed more rapidly and with fewer computational resources. A related object of the

invention is to provide such methods and apparatus as can be performed with less expensive equipment, off-the-shelf or otherwise.

Summary

The above objects are among those attained by the invention, which provides improvements in *reconstructive imaging*. In one aspect, those improvements pertain to generating a back-projection of a measured projection image (or other 2D representation) to 5 create an estimate of a 3D volume. Such back-projection can be applied in connection with conventional reconstructive imaging, as well as in connection with iterative reconstruction techniques of the type used in tomosynthesis.

In both applications, as well as any others in which such back-projection might be 10 applied, aspects of the invention provide for creating each plane of the reconstructed volume by performing a geometric warp of the projection data.

According to some aspects of the invention, e.g., those for use in connection with computed tomography systems where only the x-ray source moves relative to the subject and the detector position remains fixed relative to the subject, the reconstruction volume is positioned such that each plane of the volume is parallel to, and aligned with, the plane of the detector. 15 Related aspects of the invention employ a geometric warp that is a first order function separable in the x-axis and y-axis, e.g., capitalizing on the reduced complexity of the geometry resulting from such parallel alignment. This makes the computation simpler for both general-purpose computers as well as graphic processing units (GPUs). In further related aspects of the invention, such a first order warp is performed by mapping the projection point of the four corners of the 20 detector plane (or “projection plane”) into a specific plane of the reconstruction volume. When a back-projection according to the invention is applied (by way of example) to a computed tomosynthesis application where the detector remains stationary (relative to the subject), these 25 four points define a rectangle aligned to the axis of the detector. In this and other such applications with such a resulting alignment, aspects of the invention call for generating the step function of the resampling function (i.e., the warp) as the ratio of length of the rectangle (on each the x-axis and the y-axis) and the corresponding length of the detector (in the same axis).

Other aspects of the invention likewise provide for using a set of complimentary warps to 30 perform the back-projection, even where the projections are produced by non-aligned detectors. In these cases, the warp can be referred to as a *perspective warp* and is computationally more expensive.

A related aspect of the invention provides improvements with respect to calculating the forward-projection of a 3D volume. In this case a 3D model (for example, an estimate of the volume being reconstructed) produces a 2D projection lying in the detector plane. This projection may represent an estimate of a projection originally acquired at that focus. According
5 to this aspect, for each slice in the 3D representation parallel to the detector plane, the following steps are performed in connection with the forward projection: (i) warping the respective slice of the 3D representation to generate a respective slice of second 3D representation by applying to the former (i.e., the slice of the first 3D representation) a selected linear mapping, where that selected linear mapping would map a region defined by projection of corners of the respective
10 projection onto the respective slice of the first 3D representation for the respective focus to match dimensions of the projection in the detector plane, and (ii) incrementing values of each of one or more pixels of the 2D representation by an amount that is a function of a value of a correspondingly indexed pixel of the respective slice of the second 3D representation.

Further related aspects of the invention provide methods as described above in which at
15 least the warping steps are executed on graphics processing unit (GPU) coprocessor, e.g., of the type having programmable pixel and vertex shaders. In instances where the GPU provides an *add* instruction (or related arithmetic instructions) that provide for parallel processing of multi-component operands, further related aspects of the invention contemplate executing such an instruction with a lower-order portion of a mantissa in one component of an operand and a
20 higher-order portion of that mantissa in another component of that same operand. Carry-over (from the lower-order to higher-order portions) and/or recombination of the mantissa portions can be performed subsequently.

Further aspects of the invention provide improved methods as discussed above for use in medical imaging. According to these aspects, both the forward-projecting and the back-projecting innovations can be used as part of a maximum likelihood estimation maximization (MLEM) or in other iterative techniques.
25

These and other aspects of the invention are evident in the drawings and in the description that follows.

Brief Description of the Drawings

A more complete understanding of the invention may be attained by reference to the drawings, in which:

Figure 1 depicts a mammography system according to one practice of the invention;

5 Figure 2 depicts generation of a measured projection image in a mammography system of the type show in Figure 1;

Figure 3 depicts basic steps of an maximum likelihood estimation maximization (MLEM) reconstruction;

10 Figure 4 depicts steps performed in a forward-projection operation according to one practice of the invention, and

Figure 5 depicts steps performed in a back-projection operation according to one practice of the invention.

Detailed Description of the Illustrated Embodiment

The invention provides improved methodologies and apparatus for forward-projection and back-projection that reduce the time and processing power required to reconstruct three-dimensional volumetric representations, e.g., from a series of two-dimensional measured projection images.

The invention has application, for example, in medical imaging and the embodiment described below is one such application: a tomosynthesis mammography system, which is based on maximum likelihood estimation maximization (MLEM) algorithm. It will be appreciated that the improved forward-projection and back-projection methodologies and apparatus described herein have application in computed tomography (CT), positron emission tomography (PET), single photon emission computed tomography (SPECT), and other medical imaging applications, as well as in reconstruction applications outside the medical realm, whether based on MLEM or otherwise.

Turning to the illustrated embodiment, Figure 1 depicts a mammography system 10 according to one practice of the invention. The system includes a conventional image acquisition apparatus 12 that generates projection images 14 of the breast of a patient, depicted here partially obscured behind privacy partition 16. The projections are produced by illuminating the breast with a radiation source 22 and detected by a charge-coupled device or other 2D sensor array 20, emissions not absorbed by the breast. Multiple projections are required for adequate reconstruction of a three-dimensional representation of the breast. In one embodiment, those projections are generated in accord with the principles of computed tomography (CT), i.e., with the source 22 at discrete foci on an arc 24 that completely, or more completely, surrounds the breast. In another embodiment, those projections are generated in accord with the principles of computed tomosynthesis, i.e., with the source 22 at discrete foci along a smaller arc above the breast. Regardless, in the drawing, only three exemplary foci are shown, labeled 24a – 24c, though in actual implementation many more foci would likely be used.

In the illustrated embodiment, the image acquisition apparatus 12 comprises an x-ray source and 2D sensor of the type provided, for example, in a digital mammography system of the type available in the marketplace. Of course, the apparatus 12 may instead comprise an analog mammography system, a CT system or other imaging system, e.g., so long as the measured

projection images produced thereby are digitized or are otherwise amenable for processing in accord with the teachings hereof.

Digital data processor 26 analyzes the two-dimensional projections 14 generated by apparatus 12 to reconstruct a three-dimensional volumetric image of the breast. That processor 5 26 may form part of a workstation that performs additional functions (such as controlling acquisition apparatus 12, storing patient data, generating user interfaces, or otherwise) or it may be a stand-alone or embedded processor that works within, or alongside, such a workstation (or other computing device) to perform such analysis.

As shown in the drawing, the digital data processor includes a central processing unit 10 (CPU) 30, dynamic memory (RAM) 32, and I/O section 34, all of the type conventionally known the art. The digital data processor 26 may be coupled, via I/O section 34, with acquisition device 12, as well as with a monitor or other presentation device 28 on which the three-dimensional reconstruction is displayed.

Illustrated digital data processor 26 also includes a graphical processing unit 36 that is 15 coupled to the CPU 30, through which it can access the other elements of the digital data processor 26, as shown. The GPU 36 serves, in the illustrated embodiment, as a coprocessor, operating under the control of the CPU 30 to generate front- and back-projections from the projection images 14 in accord with the teachings hereof. Other embodiments of the invention 20 employ multiple GPUs for this purpose, each responsible for a respective portion of the reconstruction processing. Still other embodiments use no GPU at all, relying on the CPU 30 and/or other co-processing functionality (such as floating point units, array processors, and so forth) to provide or supplement such reconstruction processing, all in accord with the teachings 25 hereof. In embodiments which use GPUs, preferred such devices are of the variety having programmable vertex shaders and programmable pixel shaders and are commercially available from ATI Research (for example, the Radeon™ 9700 processor), NVIDIA (for example, the GeForce™ FX and Quadro® processors).

Figure 2 depicts generation of a measured projection image by apparatus 12 when the 30 source 22 is positioned at one focus (e.g., of a computed tomography apparatus, a computed tomosynthesis apparatus, or otherwise). X-rays emitted by the source travel through the imaged volume 39 and impinge on the sensor array 20 in a *cone beam*. A *cone beam* is defined as a collection of rays that diverged from a local point source. Those rays are detected by cells in the 2D sensor array 20 and output, typically, as pixels whose values correspond to the incident ray

intensity. Rays passing through tissue in the breast are absorbed to an extent that depends on the volume and relative density of the intervening tissue. This results in *transmission or emission shadow* 39 of varied intensity at the sensor array 20. This transmission or emission shadow is a classic 2D x-ray image. It is that intensity variation, recorded in the pixel values of the 2D sensor array 20 for taken at each x-ray source focus, which is analyzed by digital data processor 26 to reconstruct the imaged volume 39, which is drawn here with its constituent volumetric elements, or voxels.

Systems according to the invention utilize digital data processor 26 to perform that reconstruction using improved back-projections and/or forward-projections in which projection operations are performed as if rays of the source 22 had passed through the imaged volume and impinged on the sensor array — not as a divergent cone beam — but, rather, at a constant angle relative to the plane of the sensor array (i.e., the “detector plane”). Put another way, those operations are performed as if the rays had traveled along parallel paths through the volume and to the detector plane, e.g., as if emitted by a point source disposed infinitely far from the sensor. Those projection operations are executed, in the illustrated embodiment, by mapping and re-mapping coordinates of voxels and pixels (in the space defined by volume 39 and region defined by sensor array 20, respectively) prior to processing-intensive phases (i.e., so-called “inner loops”) of the reconstruction. The projection operations are preferably performed using a GPU (or GPUs) of the type referred to above, which can perform the mapping and re-mapping operations at maximal speed and with minimal code by executing first-order “warping” operations.

In the illustrated embodiment, the improved methodologies and apparatus for back-projection and forward-projection are used with an MLEM reconstruction technique; although, in other embodiments, those improved methodologies and apparatus are utilized in connection with other iterative techniques, such as the so-called Algebraic Reconstruction Technique, or non-iterative techniques. The basic steps of that MLEM reconstruction are shown in Figure 3. There, MLEM reconstruction is shown to be an iterative process involving successive stages of forward-projection (or re-projection), error-measurement, back-projection and correction. With each iteration, absorption estimates for the volume 39 converge toward a suitable representation of the imaged breast — typically, within ten to twenty iterations, though, this depends on the number of measured projection images 14, the resolution of the imaging apparatus, and the specifics of the MLEM algorithm employed.

With specific reference to Figure 3, each iteration begins with a forward-projection (or re-projection); *see*, step 40. Here, absorption estimates for the volume 39 are stored in a matrix *estimated volume*, though other data structures or storage formats known in the art can be used instead. Those estimates are integrated along each x-ray path that extends from the source 22 to each sensor pixel. This is done for each source 22 focus and results in a set of hypothetical (or synthetic) projection images that geometrically correspond to the original set of measured projection images generated by the apparatus 12. According to the prior art, the x-ray paths are treated as diverging from a local point source. However, in the illustrated embodiment (as further discussed below), the paths are treated as if parallel to one another relative to the sensor's detector plane, as if the point source 22 were infinitely far from the sensor. Prior to the first iteration, the estimated volume is initialized with an average expected absorption constant or any other default value (or set of values).

In step 42, those hypothetical reconstructed projection images are compared with the measured projection images generated by the image acquisition apparatus. For each x-ray source focus, an array (or other data structure or storage format) *difference* is generated representing pixel-by-pixel differences between intensities in the measured projection images and current hypothetical projection images for the current x-ray source focus. In embodiments that utilize a Poisson noise-based weighting in the correction-phase (step 46), an array (or other data structure or storage format) *reference* can also be generated in step 42 representing the minimum expected error assuming Poisson noise in the original projection images. The difference array — and, where used, the reference array — have the same dimensions as the measured and hypothetical projection images and, thus, are suitable for back-projection into “correction” volumes having the same dimensions as the estimated volume.

(A discussion of the aforementioned Poisson noise-based weighting is provided in Wu *et al*, “Tomographic Mammography Using A Limited Number Of Low-dose Conebeam Projection Images,” *Med. Phys.*, pp. 365, *et seq.* (2003) (and, specifically, the discussion of the use of such weighting at p. 371) and Lange *et al*, “EM Reconstruction Algorithms For Emission And Transmission Tomography,” *J. Comput. Assist. Tomogr.* 8, pp. 306, *et seq.* (1984), the teachings of both of are incorporated herein by reference.)

In step 44, a first correction volume is generated by back-projection of the difference array over all x-ray source foci. A second such volume can also be generated, in embodiments that utilize the aforementioned Poisson noise-based weighting, by back-projection of the reference array over those same x-ray source foci. As above, whereas according to the prior art,

these back-projections are performed along x-ray paths that diverged from a local point source, in the illustrated embodiment (as further discussed below) the paths are assumed parallel to one another relative to the sensor's detector plane — again, as if the source 22 were infinitely far from the sensor.

5 In step 46, the absorption estimates in the voxels of the estimated volume are updated or corrected. In some embodiments, the value of each voxel of the current estimate matrix is reduced (or increased) by the value of the corresponding voxel of the first correction volume. In embodiments that utilize the aforementioned Poisson noise-based weighting, the reduction (or increase) is by the ratio of the values in the corresponding voxels of the difference and reference arrays.

10 Steps 40 – 46 are repeated the absorption values in the estimated volume matrix converge to close approximation of the actual object, again, typically, within ten to twenty iterations, depending on the number of projection images 14, the resolution of the imaging apparatus, and the specifics of the MLEM (or other) reconstruction algorithm employed. It should be noted that 15 even after the algorithm stops converging the estimated volume matrix will retain some error terms, and that these error terms will continue to be redistribute throughout the volume between subsequent iterations of the algorithm.

20 As noted above, in prior art, forward projection and back-projection are performed along x-ray paths, that diverge from a local point source 22 as they head toward the different detector elements of the 2D sensor array 20. A problem with this is that it complicates the inner loop of each of these projections. For forward-projections, for example, it requires calculation of the intersection of each estimated ray with each voxel along the ray's path in order to determine how much that voxel's estimated absorption decreases the ray intensity. It also requires an additional 25 steps, within the inner loop, to determine neighboring voxel's contributions to absorption (this is particularly true in instances where the volume 39 is defined with voxels that are "taller" (in the z-dimension) than they are wide and/or deep (in the x- or y-dimensions). Similar problems are encountered in back-projections according to the prior art.

The illustrated embodiment overcomes this by performing forward-projection (step 40) in the manner illustrated in Figure 4.

In step 40a, the digital data processor 26 begins a loop through the foci 24a – 24c. In step 40b, the estimated projection image for the current focus is zeroed.

5 In steps 40c – 40f, the digital data processor 26 loops through the slices (i.e., z-axis indices) of the *estimated volume* array and generates corresponding slices in an array (or other data structure) *warped volume*. This is done by a first-order warp which linearly maps the voxels of each slice of the former (*estimated volume*) to the voxels of the latter (*warped volume*).

10 The result of the processing effected in steps 40c – 40f is to produce a volume of absorption estimates warped as if rays of the source 22 had passed through the imaged volume 39 — not as a diverged cone beam— but, rather, at a constant angle relative to the plane of the 2D sensor array 20 , or detector plane. Put another way, the effect is as if the rays had traveled along parallel paths, e.g., as if emitted by a point source disposed infinitely far from the sensor.

15 Once this is accomplished, the inner loop of the forward-projection can be executed as shown in steps 40g – 40j. Particularly, the value of each pixel in the hypothetical projection can be determined by summing absorption constants of correspondingly indexed voxels in the 20 respective slices of the *warped volume* array. Though not shown in the flow chart, each such sum is preferably weighted by a factor dependent on the radial distance of the hypothetical projection pixel from a point of intersection of a normal extending from the surface of the sensor array 20 to the focus for which that projection is being generated. In one embodiment, that weighting decreases from one (for pixels at that point of intersection) = L / L (where L = length of the x-ray source normal from x-ray source to the detector plan) to $SQRT(L^2 + dx^2 + dy^2) / L$ (where dx and dy are the offsets from the intersection of the x-ray normal to the indicated pixel in the detector), though, in other embodiments, this may vary.

Pseudocode corresponding to the forward-projection steps 40a – 40j is reprinted below:

```
25  for focus = 1 to (n)
      # start with blank estimated projection images
      for u = 1 to x(max)
          for v = 1 to y(max)
              estimated projection(foci,u,v) = 0
30
      # warp volume as if illuminated by a point source at
      # infinity (i.e., parallel rays)
      for k = 1 to z(max)
          for u = 1 to x(max)
```

```

        for v = 1 to y(max)
            warped_volume(u,v,k) = warp1(estimated_volume,u,v,k,
                                         location(sensor),foci)
    # perform forward-projection (re-projection)
5      # using warped volume
      for k = 1 to z(max)
          for u = 1 to x(max)
              for v = 1 to y(max)
                  estimated_projection(foci,u,v) = estimated_projection(foci,u,v)
10                 + warped_volume(u,v,k)

```

In this example, the function *warp1()* is the aforementioned first-order warp. It linearly maps a voxel at coordinates (u,v,k) of the current slice, k , of the *estimated volume* array to a voxel that lies within a volume defined by a projection, at the angle of the current focus, of the 15 corners of the sensor array 20 onto the volumetric plane of that slice. Put another way, the function can be considered as projecting the corners of the sensor array 20 onto the plane of the current slice k of the *estimated volume* matrix, given the current focus, and as linearly mapping voxels of that projection (which can be thought of as a “shadow” of the sensor) — including those of the current slice of *estimated volume* that fall within that “shadow” — so that they 20 match the dimensions of the sensor in the original sensor plane.

The illustrated embodiment overcomes the deficiencies of prior art back-projection (step 44) by executing that operation the manner illustrated in Figure 5. In step 44a, the digital data processor 26 zeros the volume correction matrices into which the difference and reference arrays will be back-projected and which, in step 46, will be used to correct the estimated volume 25 matrix. Step 44b is the start of a loop through individual slices of the volume, while steps 44c – 44d begin loops through the individual voxels of that slice. In steps 44e – 44g, the digital data processor loops through the foci 24a – 24c and calculates, for each, a pixel of a warped difference array for the current slice and focus, and a pixel of a warped reference array, again, for the current slice and focus.

30 The warped difference array is generated from corresponding pixels of the difference array generated in step 42. This is done by a first-order warp which linearly maps the pixels of the difference array.

35 In embodiments which utilize a second volume correction matrix in the correction step 46, the digital data processor generates a warped reference array in step 46g. This is accomplished in the same manner as the warped difference array.

The result of the processing effected in steps 44f – 44g is to produce an array of intensity differences (and, where applicable, an array of reference intensity values) warped as if rays of the source 22 had passed through the imaged volume 39 — not as a collection of rays that diverged from a local point source— but, rather, at a constant angle relative to the plane of the 5 sensors detector plane 20. Put another way, the effect is as if the rays had traveled along parallel paths, e.g., as if emitted by a point source disposed infinitely far from the sensor.

Once this is accomplished, the inner loop of the back-projection can be executed as shown in steps 44h – 44i. Particularly, the value of each voxel of the current slice of the first 10 volume correction matrix is determined by summing intensities from correspondingly indexed pixel values of the warped difference array. And, likewise, in embodiments that utilize a second volume correction matrix, the value of each voxel in the current slice of that matrix is determined by summing intensities from correspondingly indexed pixel values of the warped reference array.

Pseudocode corresponding to the back-projection steps 44a – 44i is reprinted below:

```

15      # start with blank volume correction matrices
      for i = 1 to x(max)
          for j = 1 to y(max)
              for k = 1 to z(max)
                  a(i,j,k) = 0
                  b(i,j,k) = 0
20
      # back-projection loop
      for k = 1 to z(max)
          for j = 1 to y(max)
              for i = 1 to x(max)
                  for focus = 1 to (n)
                      # warp difference and reference arrays as if
                      # illuminated by a point source at infinity
                      # (i.e., parallel rays)
                      warped_difference(i,j,k) = warp2(difference,i,j,
                                              k,location(sensor),foci)
30
                      warped_reference(i,j,k) = warp2(reference,i,j,k,
                                              location(sensor),foci)
                      # back project error terms into volume correction
                      # matrices a & b
35
                      a(i,j,k) = a(i,j,k) + warped_difference(foci,i,j)
                      b(i,j,k) = b(i,j,k) + warped_reference(foci,i,j)

```

In this example, the function *warp2()* linearly maps the pixels of the difference array (and reference array) to pixels of the warped difference array (and warped reference array). This 40 is effectively the complement to *warp1()*, discussed above, and hence might be thought of as an “unwarp.” It linearly maps pixels at coordinates (i, j) for the current focus of the *difference* array

to a pixel that lies within a region defined by a projection, at the angle of the current focus, of the corners of the current slice k of the estimated volume onto the plane of the sensor array 20. Put another way, the function can be considered as projecting the corners of the current slice of the estimated volume onto the plane of the sensor array 20, given the current focus, and as linearly 5 mapping pixels in that region (which can be thought of as a “shadow” of the slice) — including those of the *difference* array (and *reference* array) that fall within that “shadow” — so that they match the dimensions of the slice in its actual volumetric plane.

In the illustrated embodiment, the forward-projection and back-projection operations (depicted in steps 40a – 40j and 44a – 44i, respectively), if not the entire volumetric 10 reconstruction (depicted in steps 40 – 46) are executed using a GPU (or GPUs) of the type referred to above. These perform the linear mappings that are the basis of the *warp1()* and *warp2()* functions as first-order warps, which can be executed rapidly and with minimal code.

To accommodate the limited floating point precision of the GPUs currently available in the marketplace, the illustrated embodiment capitalizes on another aspect of their native 15 architectures to attain the necessary precision for certain addition operations. Those architectures assume that many operands have four components, namely, three RGB color values and one intensity value, which are processed in parallel. Unfortunately, the precision of each component is typically limited to n bits, where n is typically 24 bits, or less.

For addition operations requiring greater than n bits of accuracy, the illustrated 20 embodiment uses two (or more) of these limited precision components to represent the lower- and higher-order bits of the operands’ mantissas. Thus, for example, for additions requiring 32-bit accuracy, the low-order 16-bits of the mantissas of each operand can be retained and processed stored in one of the components (say, the Red component), while the most significant 16-bits are retained and processed via another of the components (say, the Green component). 25 Of course, a different bit-wise dividing point can be used. Following the GPU’s execution of an *add* instruction on such operands, the lower-order component of the instruction’s target is tested for overrun and the high-order component incremented with a carry, as necessary. In other embodiments, a similar use of multi-part mantissas is used in connection with other arithmetic operations, such as *subtract*; *complement*, and so forth. Regardless, prior to executing a *multiply*, 30 or other instruction, these multi-component mantissas are recombined into a single component, with the attendant loss in accuracy.

Following execution of the volumetric reconstruction and, particularly, forward- and back-projections detailed above, the digital data processor outputs the *estimated volume* matrix, e.g., for display on device 28, or for further processing by other apparatus and methods required for the attendant application.

5 Described herein are methods and apparatus achieving the objects set forth above. It will be appreciated that the illustrated embodiment is merely an example of the invention and that other embodiments incorporating changes therein fall within the scope of the invention. For example, it will be appreciated that the forward-projection and/or back-projection operations detailed herein have application in reconstructions performed using techniques other than

10 MLEM, whether iterative or otherwise, and that they have application outside the field of medical imaging. It will also be appreciated that these forward-projection and back-projection operations have application in imaging objects other than the breast. Further, it will be appreciated that the invention covers forward-projection and back-projection operations incorporating mapping and other variations from those detailed above. In view of these and

15 other modifications, what we claim is: